



Status of the Scientific Computing Program at the Laboratory

Elizabeth Sexton-Kennedy

Fermilab PAC

29 Jun 2020

Outline

- Response to January 2020 PAC recommendations
- Updates on priority projects
- AI/ML

Recommendation 1: Scientific Computing Planning

- *SCD should present a detailed report to the PAC on how it plans to satisfy the very diverse computing needs of the various experiments.*
- This plan should present the currently available hardware for all experiments and projects, and the users' usage patterns (such as how often the resources are used by experiments who have not contributed hardware).
- This report should also have information about the current support model for the users and the hardware, and point out where SCD anticipates major challenges in the future and how these will be addressed

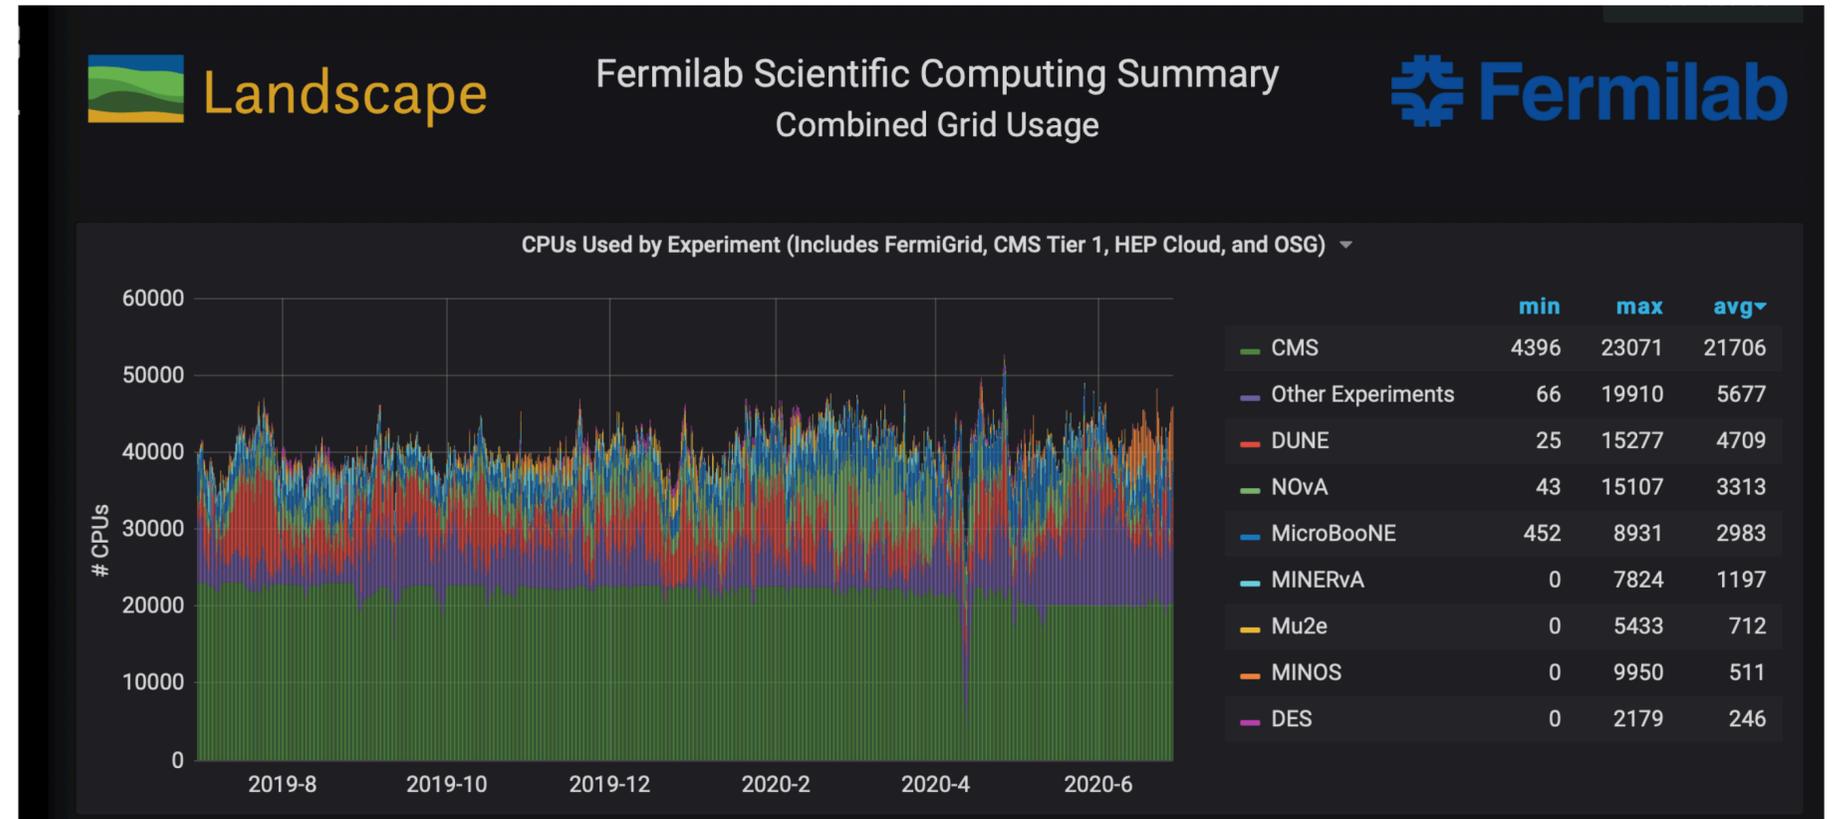
Report Highlights

- The PAC response report was made available this past weekend.

CMS scrutinized by the CERN-CRSG, funded through USCMS Fermilab experiments moving to that model, funded through IF detector and computing Ops budget.

Current FNAL HW Capacity

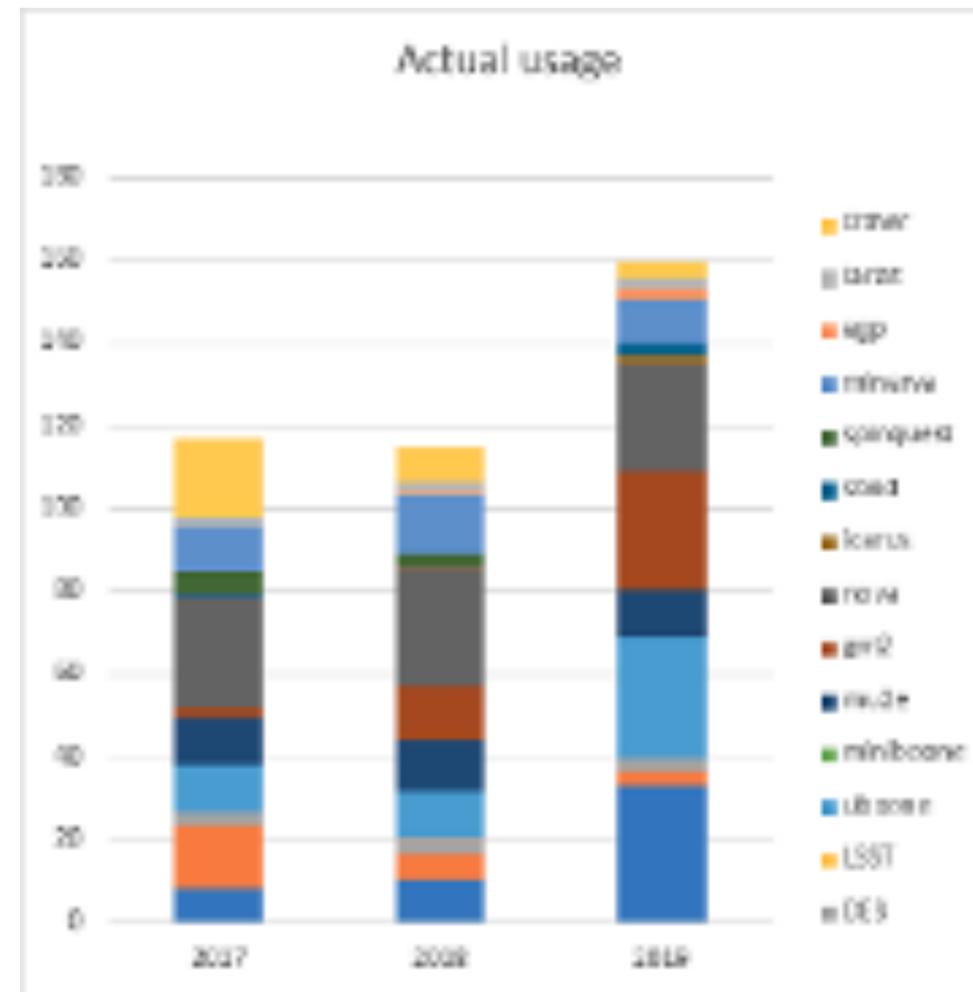
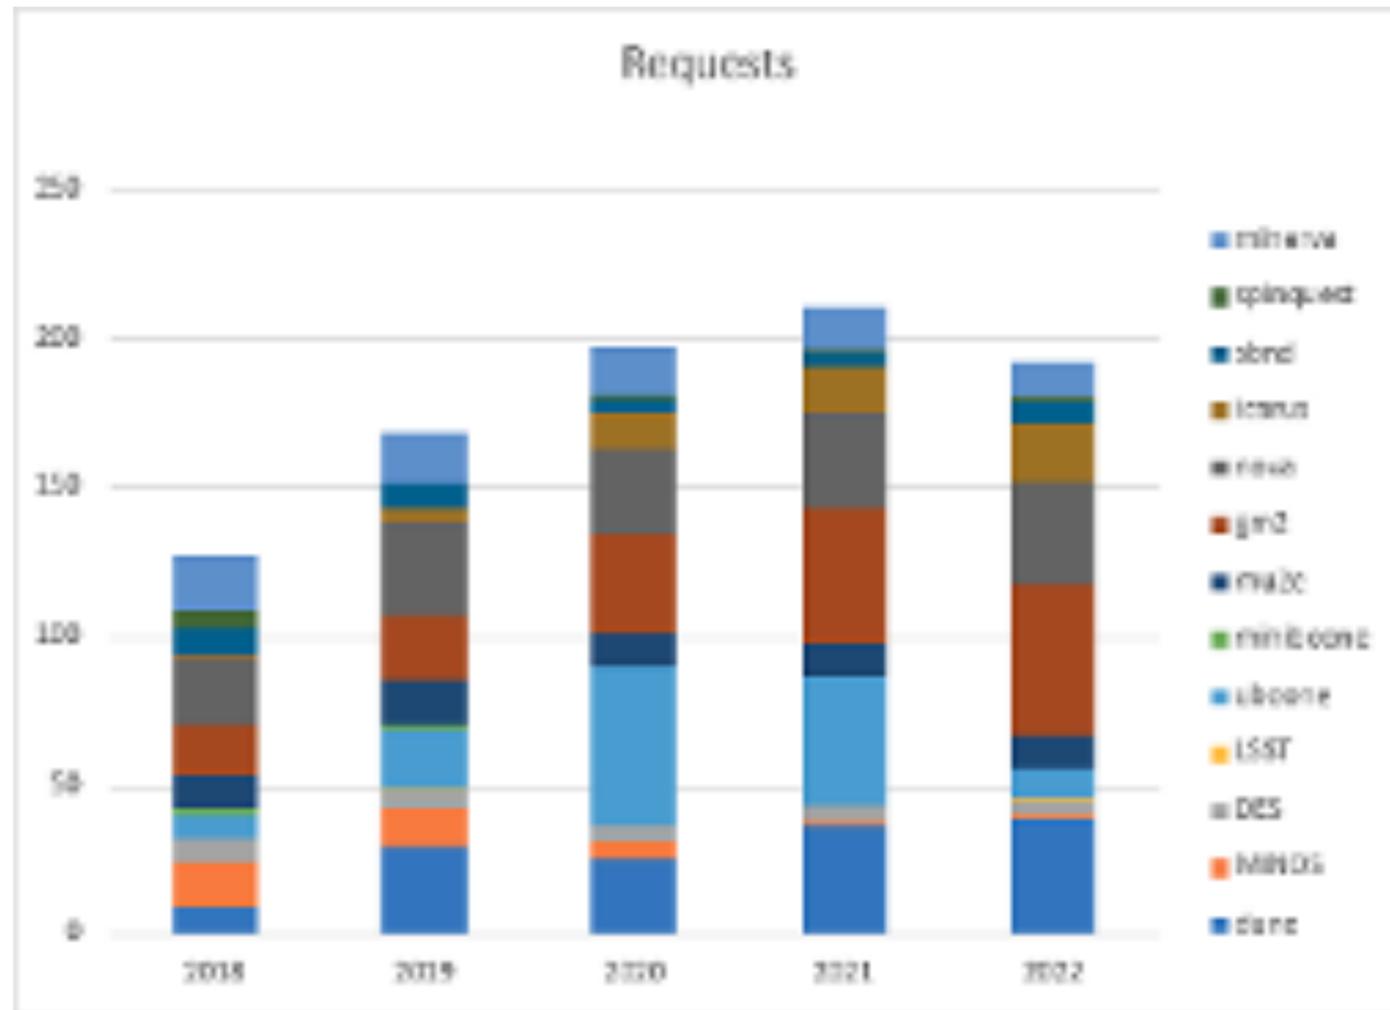
	CPU annual hours	Disk	Tape
CMS	236M	33PB	68PB
Intensity and Cosmic Frontier	210M	15PB	109PB



- Have moved to a service model.
- Funding has kept resources for both frontiers about equal.

Historic CPU Resource Trends

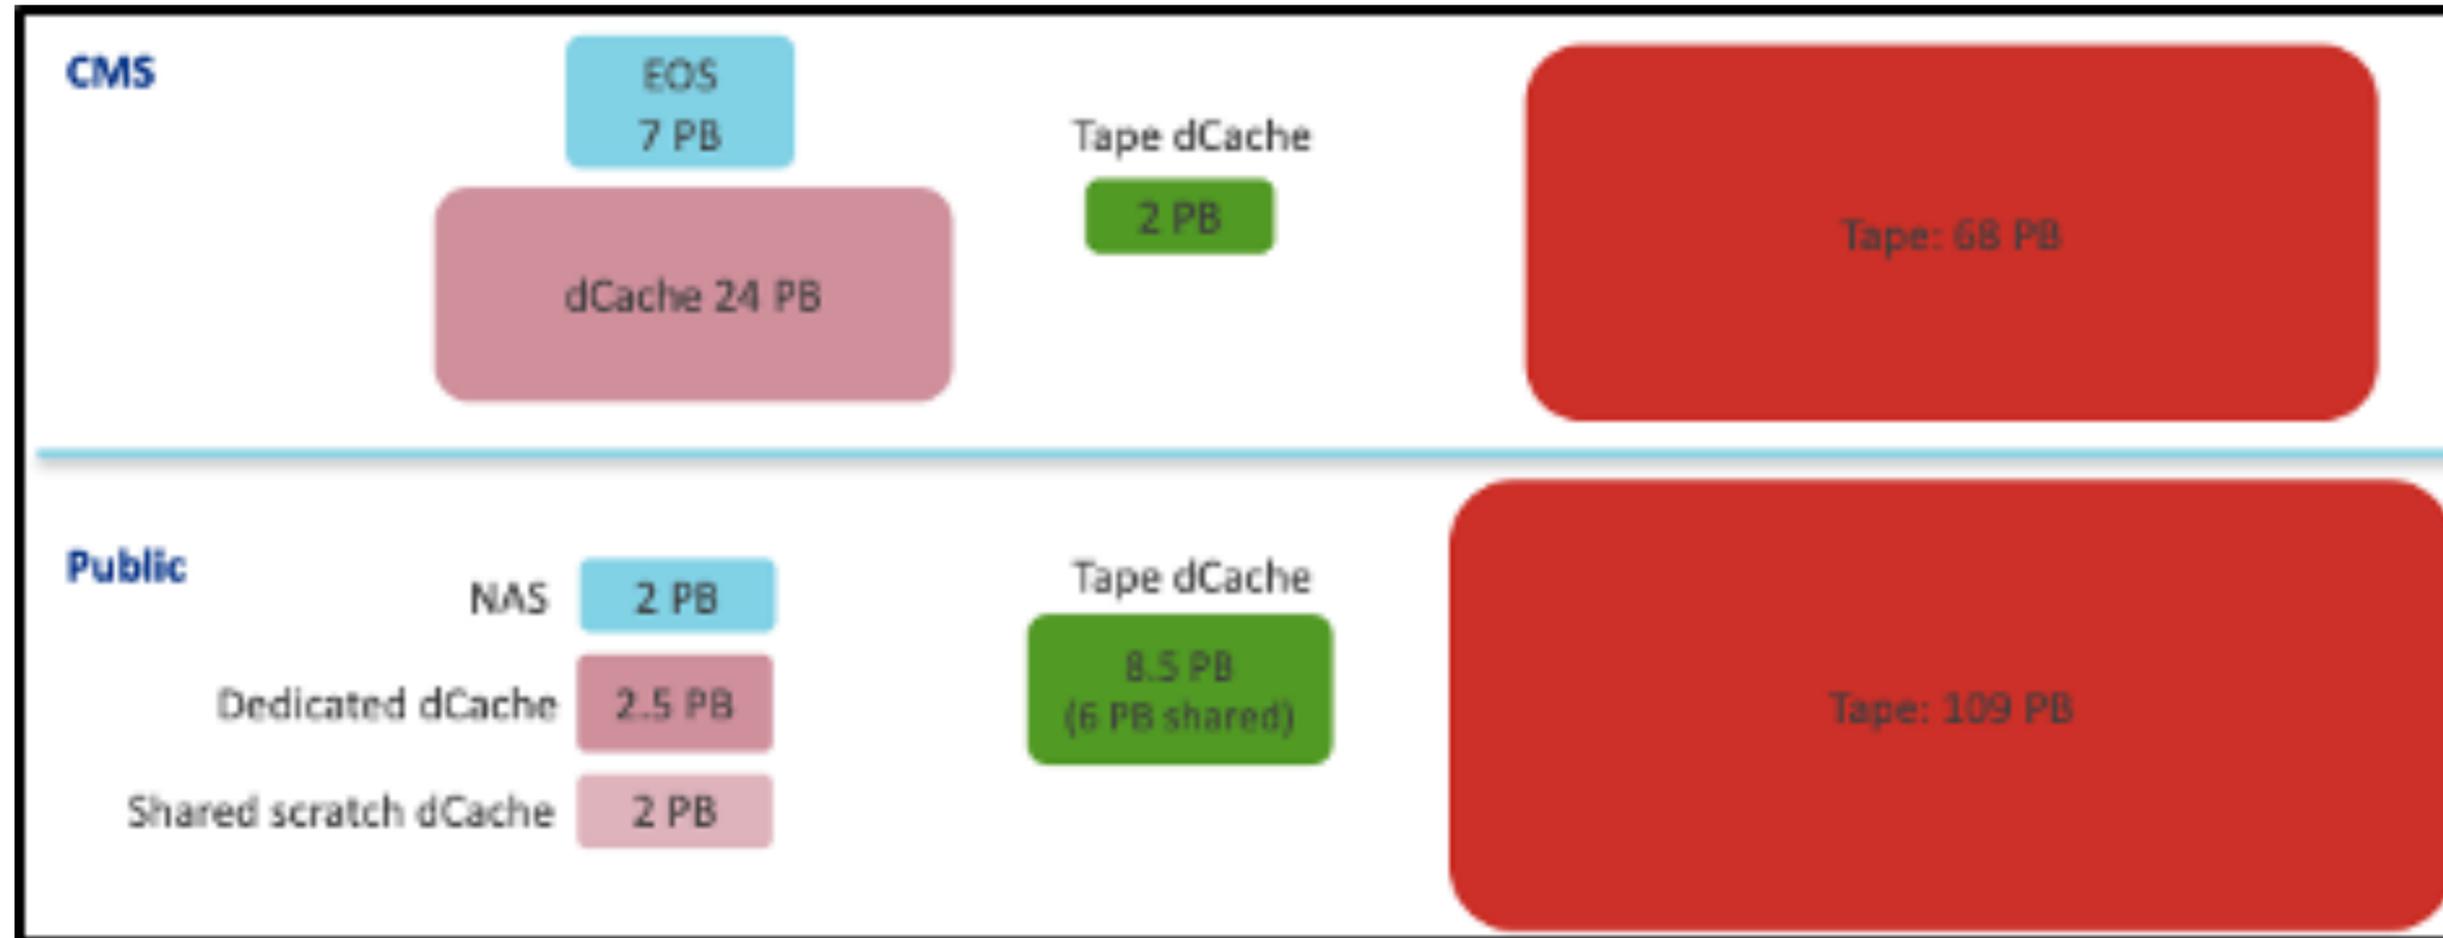
- At a capacity of 210M hours/year, the Fermilab batch processing farm meets the current needs of the IF community.
- CMS fully utilizes its 236M hours/year.



In addition
Fermilab
experiments used:
* 144M CPU
hours at HPC
Centers
* 33M CPU
hours on OSG

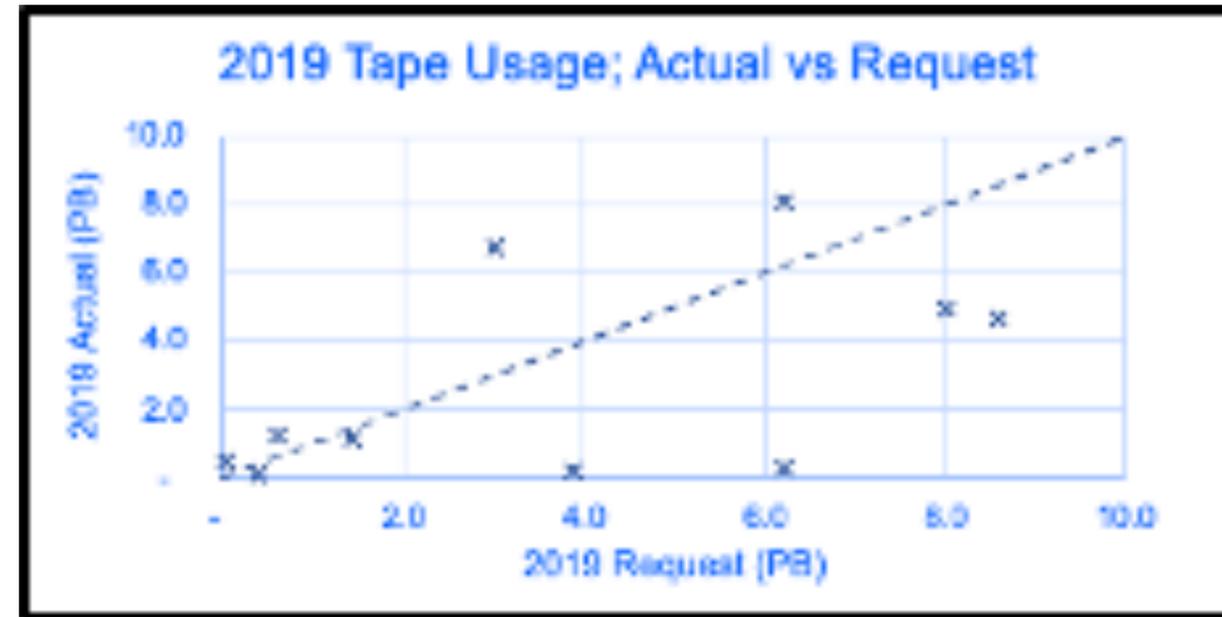
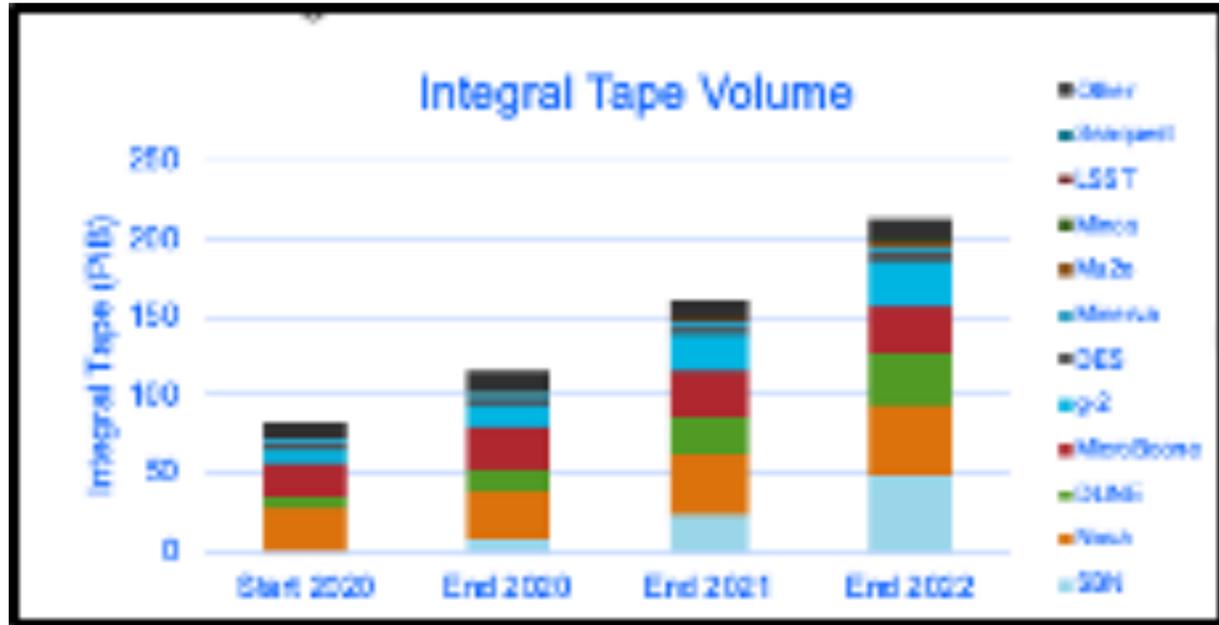
Historic Storage Architecture Differences

- CMS uses EOS for analysis disk, dCache and Tape-dCache managed separately allowing more efficient use of tape system by production system.
- IF still uses their storage as an HSM. Aim to make IF more like CMS.

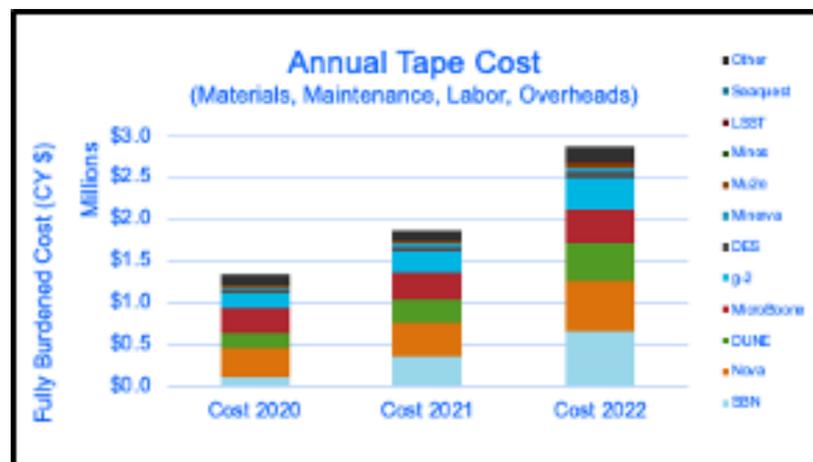


Historic Storage Resource Trends

- Historical Accuracy of Tape Requests



- Annual Tape Cost



- <my opinion>This growing trend of increased dependence on tape is in the wrong direction. dCache operations is already strained enough.<\my opinion>
- 1.4M\$ -> 2.8M\$ in 2 years is concerning.

Funding Requests and Priorities for FY21

IF Computing Thrust Walk-down

Walk Down	Description	Thrust	FTEs	FTEs	FY21
Priority	Optimal		Existing	New	25,000,000
4	Drop GPU cluster acquisition	Computing			500,000
3	Drop CPU replacements	Computing			500,000
2	Reduce experiment support	Computing	2.6		1,000,000
1	Reduce developer support	Computing	3.1		1,200,000
	Target FY21 Base				22,800,000

EF Computing Thrust Walk-down

Walk Down	Description	Thrust	FTEs	FTEs	FY21
Priority	Optimal		Existing	New	\$28,699K
1	Planned refurbishments and replacement of end-of-life CPU hardware				\$949K
2	Planned disk storage system refurbishments and buy-ahead for Run 3 capacity needs				\$1,750K
	Target FY20 Base				\$26,000K

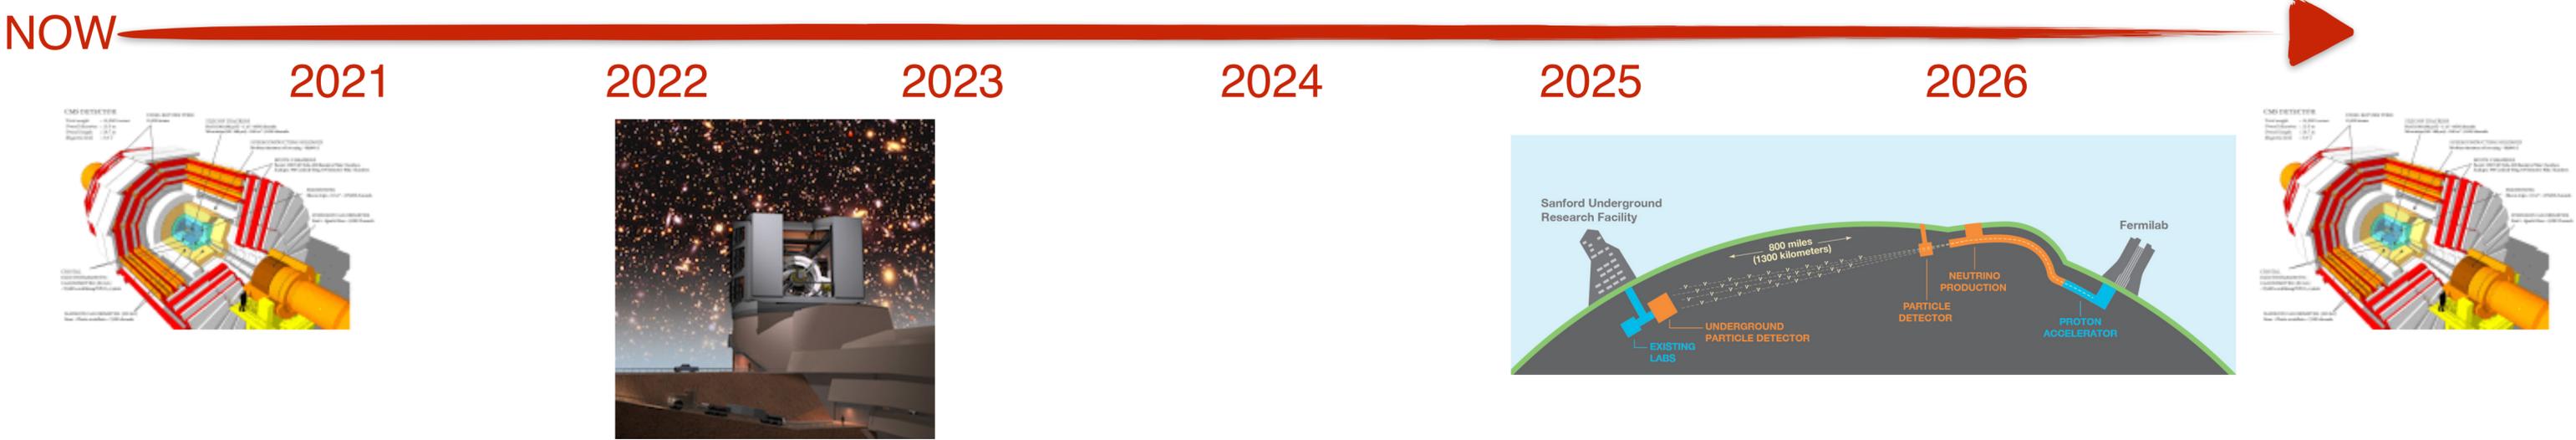
- Just finalize our request to DOE last week.
 - CMS requests include a bump in equipment spending for Run 3
 - IF includes computing investments for equipment refresh and modernization
 - Further reductions in personnel would be extremely challenging

Legacy Survey of Space and Time

LSST is the Elephant in the room

Fermilab will propose to host the next big cosmic frontier data facility. Plan to combine with NCSA once the FOA is announced. Our service model is a strength here.

Timing is favorable:



Strategy: Capitalize on our experienced facility staff

New hires come in at junior levels to become the experts of the future. We have already seen economies of scale in EF + IF staffing.



Planning Exercises

- Both Mu2e and DUNE pre-Ops planning documents were delivered to DOE with significant participation from SCD.
- CMS' HL-LHC R&D planning was delayed due to COVID-19 still on track for Sep. delivery.
- Future networking needs; the “HEP-ESNET 2020 Requirements Review” is the next important joint planning exercise. Overseeing needs for 3 groups:
 - CMS
 - Muon program - g-2 and Mu2e
 - Neutrino program - SBN, protoDUNE, DUNE
- Snowmass
 - Oliver Gutsche is one of 3 conveners of the Computational Frontier
 - Giuseppe Cerati and Daniel Elvira are topical group conveners

Recommendation 2: Prioritize Software R&D

- *The PAC requests that SCD prepare a ranked list of high priority software R&D projects (including "sustaining capabilities")*
- in consultation with the HEP community, explaining both the impact on projected future resource challenges, and how these high-priority developments can be sustained over the long term in *realistic funding* scenarios.

Optimistic Funding - Software and Facilities R&D

			FY19 Actual	FY20 Budget	FY21 Request	FY22 Request
KA 24 01 022 COMPUTATIONAL HEP - SCIENTIFIC COMPUTING	ComputationalHEP	SWF (\$K)	\$2,240	\$1,651	\$2,357	\$2,346
		M&S (\$K)	\$208	\$46	\$25	\$25
		**SUM DOLLARS (\$K)	\$2,448	\$1,697	\$2,382	\$2,371
		Scientists	1.3	1.8	1.7	1.6
		Post-docs	0.4	1.6	1.5	1.4
		Total Other FTE	6.3	3.9	5.4	5.3
		FTE	8.1	7.3	8.6	8.2
KA 24 01 045 Computational HEP - AI & ML	ComputationalHEP	SWF (\$K)		\$625	\$1,010	\$1,004
		M&S (\$K)		\$0	\$0	\$0
		**SUM DOLLARS (\$K)		\$625	\$1,010	\$1,004
		Scientists		1.0	2.0	1.9
		Post-docs		0.0	0.0	0.0
		Total Other FTE		0.0	0.0	0.0
		FTE		1.0	2.0	1.9
Grand Total		SWF (\$K)	\$2,240	\$2,276	\$3,366	\$3,349
		M&S (\$K)	\$208	\$46	\$25	\$25
		**SUM DOLLARS (\$K)	\$2,448	\$2,322	\$3,391	\$3,374
		Scientists	1.3	2.8	3.7	3.5
		Post-docs	0.4	1.6	1.5	1.4
		Total Other FTE	6.3	3.9	5.4	5.3
		FTE	8.1	8.3	10.6	10.1

- 2020 budget partially restore in the 2nd quarter.
- Maintained efforts by shifting efforts by computing professionals to scientist and postdocs.
- Requesting restoration of funding in 2021.

SCD R&D Priorities

- SCD's R&D strategy is to capitalize on our traditional strengths in areas identified by the S&C R&D community white paper but not covered by NSF. They are:
 1. Highest priority is Storage R&D
 - The focus in FY21 is on the data lake prototype project to be performed in collaboration with BNL and our international partners.
 - The data lake concept is a promising way of providing a federated system for the optimized use of various storage types with different QoS.
 2. Modernization of physics codes
 - CCE - portability frameworks and fine grain I/O
 - Creating a GPU enabled Geant application.
 - Celeritas is high risk and high reward R&D

Center for Computing Excellence Launched

- Focus is on enabling HEP application software to run effectively on DOE HPC facilities.
- Re-engineering HEP codes to use GPUs is the primary goal of the CCE.
- Portable Parallelization strategies started in January.
- Fine-Grain I/O and Storage started in March.

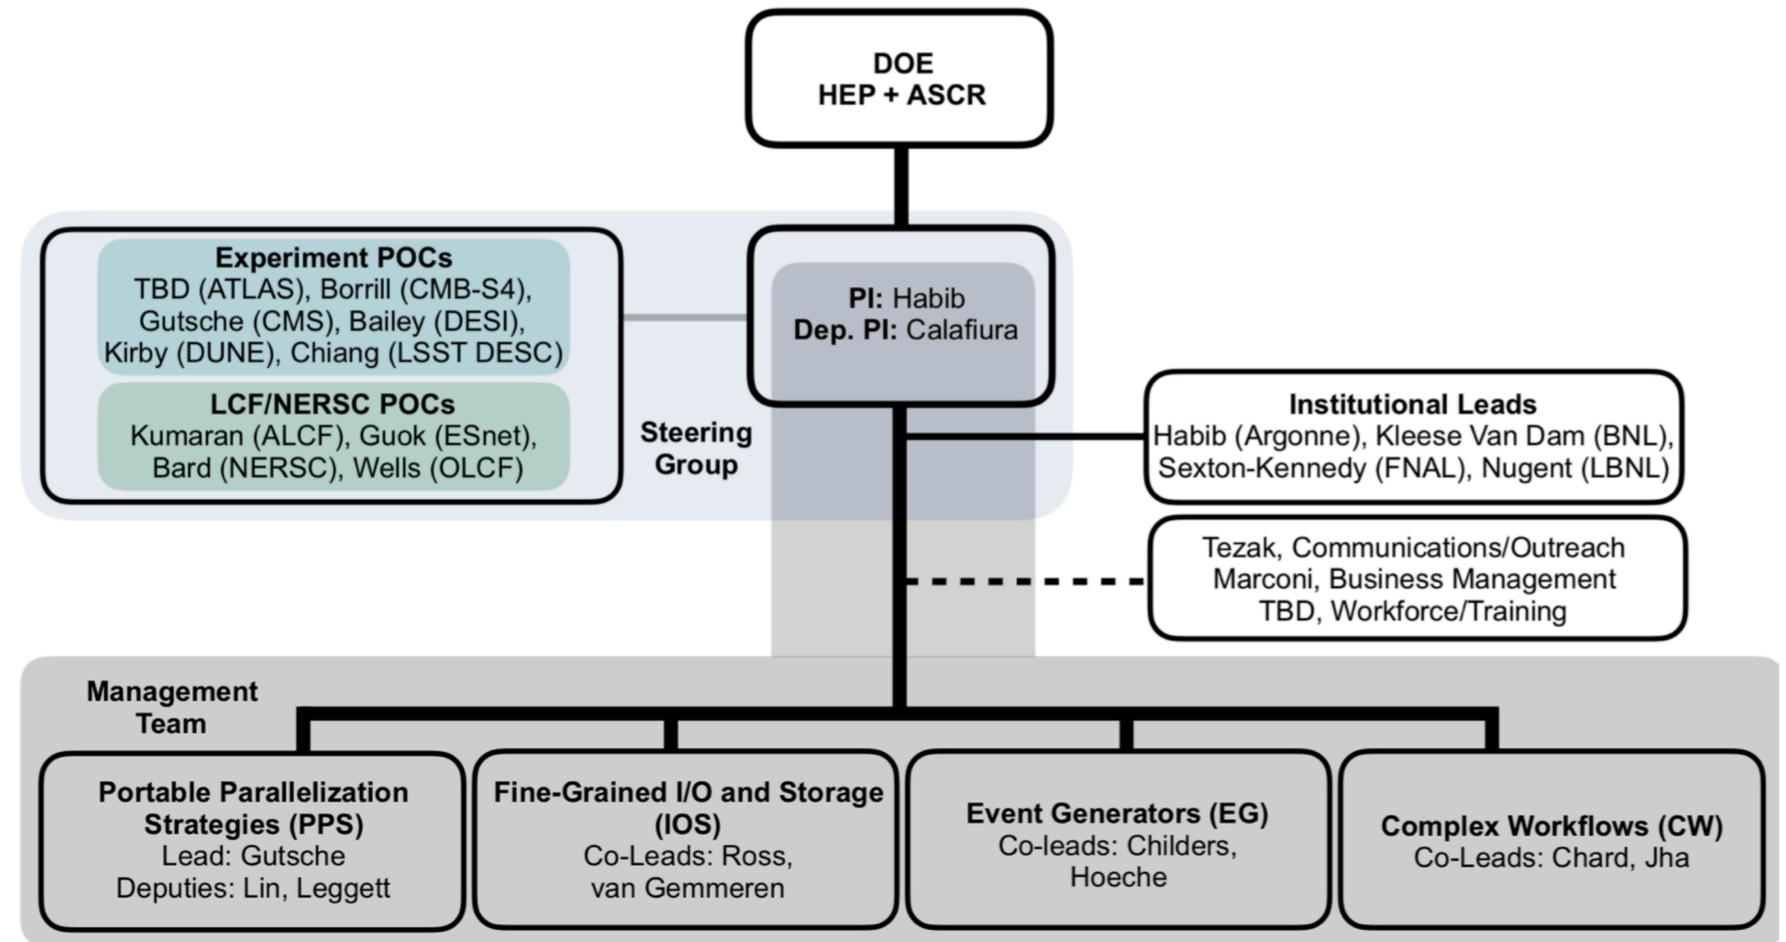
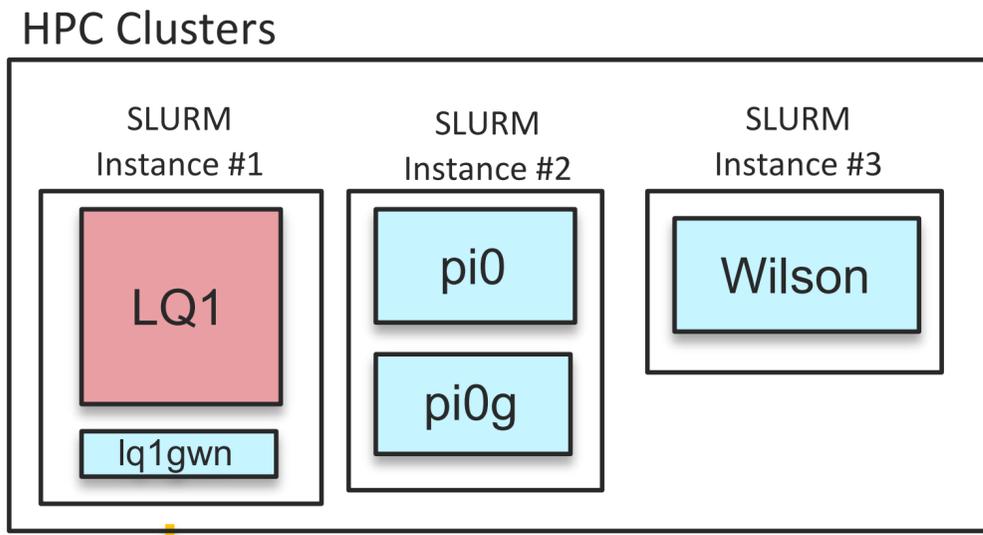
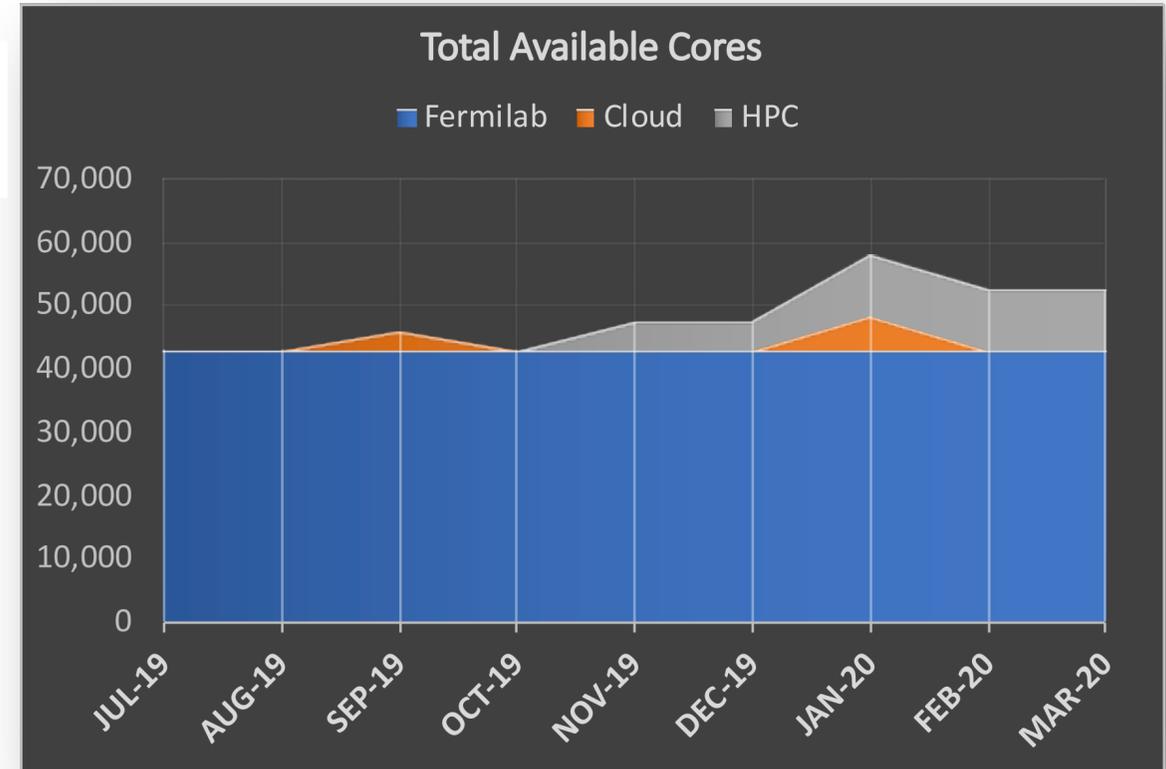


Figure 1. *Project organizational structure: The Management team is responsible for program execution, the Steering group provides guidance and feedback, and Institutional leads are responsible for technical and human resources at their institutions, as well as working together on cross-Institution interactions.*

Artificial Intelligence and Machine Learning

- Interest in AI/ML is creating funding opportunities that we must capitalize on. Recent thread started with Cheryl Ingstad the new director of the Artificial Intelligence and Technology Office of DOE is encouraging.
- Requires HEP scientific effort paired with AI practitioners.
- Will help us with our HPC challenges. There are many cross-cutting applications already developed for HEP. These need to be extended.
 - Future computing platforms will be designed to accelerate AI/ML applications =>
 - Recast algorithmic codes into ML workflows

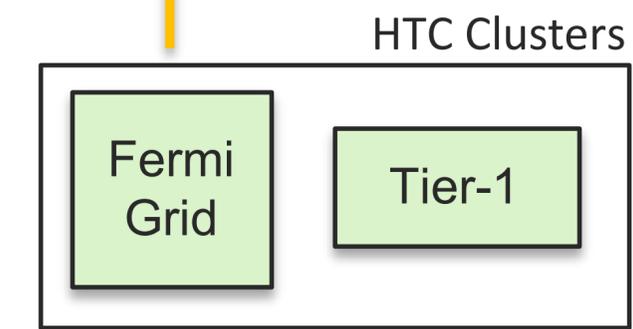
Wilson Memorial Institutional Cluster



1. Need to have CE in front of LQ1+Wilson.
2. HEPCloud project develop Decision Engine.

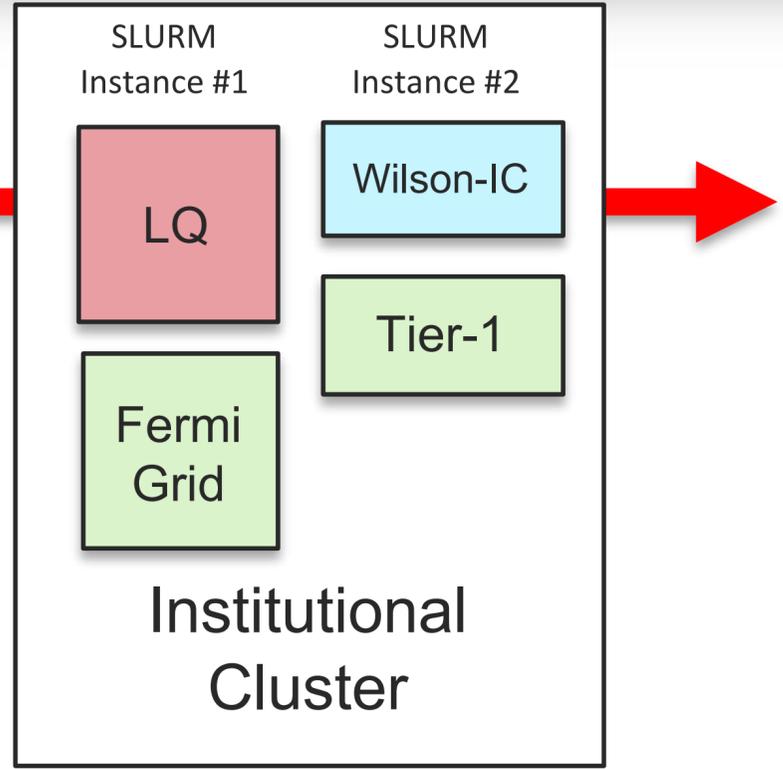
NOW

FUTURE



Create a friendly HPC Cluster w/ outbound networking

HEPCloud project develop Decision Engine.



Other Facilities R&D

- Continued evolution of the HepCloud portal - moving beyond homogeneous cloud provisioning to the “unfriendly” heterogeneous DOE HPCs
- An Elastic Analysis Facility R&D project - modernizing our analysis facilities using standard container orchestration suites
- Federated Identity - enabling Fermilab’s scientific users access to computing resources at Fermilab as well as collaborating organizations without having to go through a separate identity vetting and account creation process for each organization.
- Federated Access - providing an authorization process and framework for deciding whether a user has authority to access the requested resources.

Summary

- SCD has made good progress in communicating with next generation experiments to understand their needs. Some planning efforts completed and some started including Snowmass.
- CPU needs can be satisfied (not as clear for storage especially tape)
- SCD is still understaffed in the areas of R&D. We have asked for more effort in the areas of storage, Geant-Celeritas, and AI/ML; one each in this years CompHEP FWP.
- We did get the planned funding for CCE. Two of the 4 efforts have been started this calendar year.

Continued Education and Workforce Development

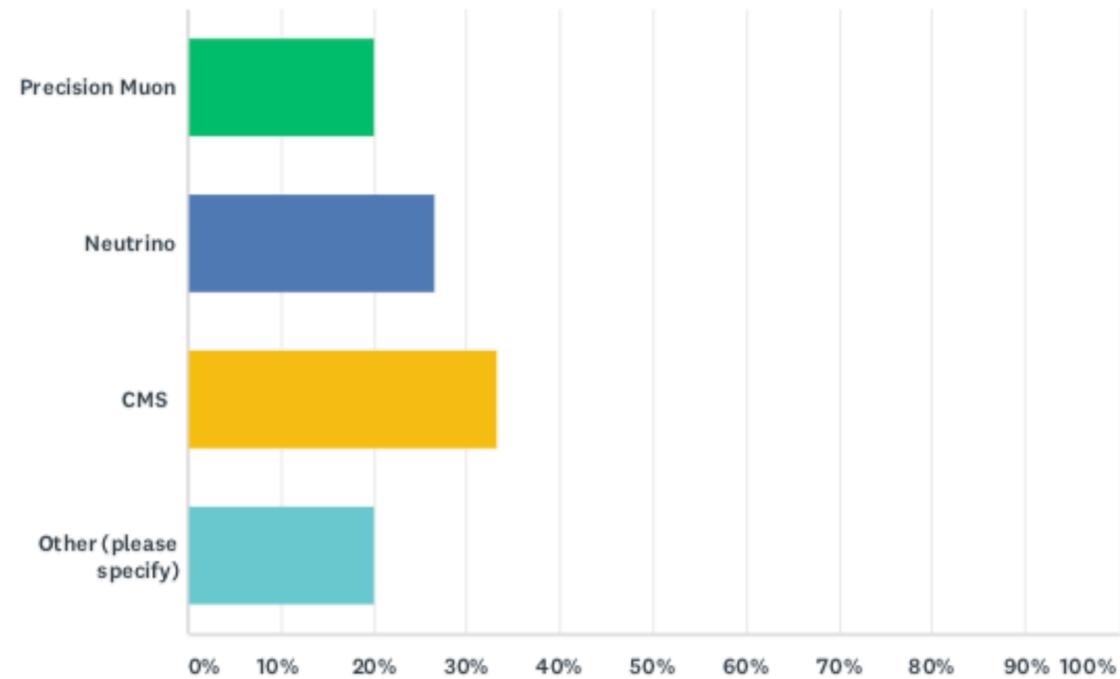
- Sponsoring PyHEP virtual workshop in July.
- Will repeat the successful C++ training course this August virtually.

C++ Training at Fermilab - Evaluations

C++ Course Evaluation Form

Q1 What part of the Fermilab program do you participate in?

Answered: 30 Skipped: 0

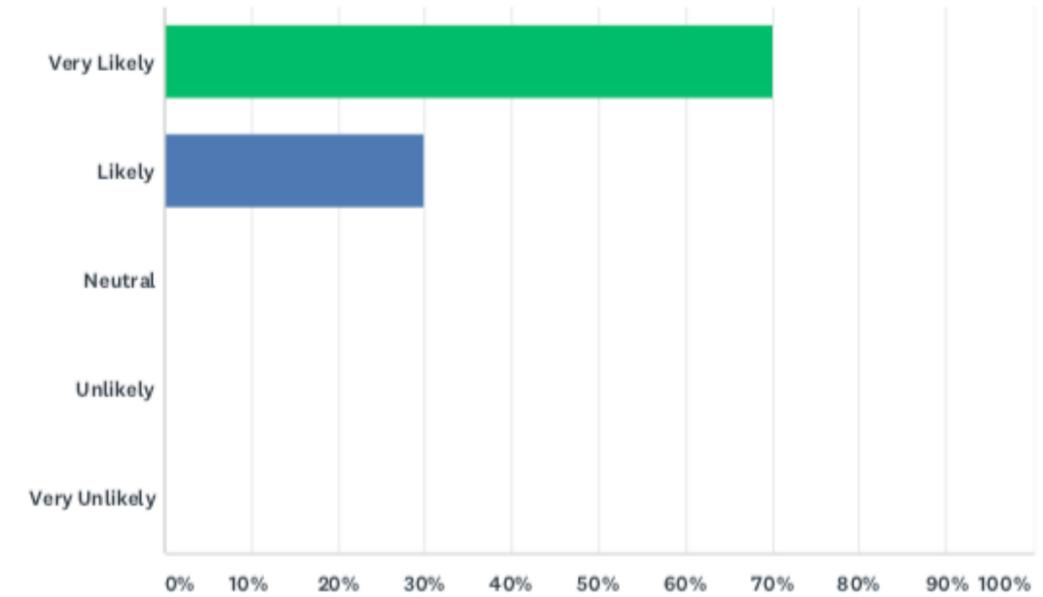


ANSWER CHOICES	RESPONSES	
Precision Muon	20.00%	6
Neutrino	26.67%	8
CMS	33.33%	10
Other (please specify)	20.00%	6
TOTAL		30

C++ Course Evaluation Form

Q9 How likely are you to recommend this workshop to a friend or colleague?

Answered: 30 Skipped: 0



ANSWER CHOICES	RESPONSES	
Very Likely	70.00%	21
Likely	30.00%	9
Neutral	0.00%	0
Unlikely	0.00%	0
Very Unlikely	0.00%	0
TOTAL		30

Lab S&C FTEs by I

- S&C funding is piece
- Even with almost (98
- Makes it harder to ab
- CompHEP was cut th
and April
- 143 is current headcc

